

Curriculum Vitae

Joseph Keshet

Personal Information

Last name	Keshet	Cell phone	+1 773 870 6481
First name	Joseph	Office phone	+1 773 834 6850
Office address	Toyota Technological Institute 6045 S. Kenwood Ave. Chicago, Illinois 60637, USA	E-mail	jkeshet@ttic.edu
		URL	www.ttic.edu/keshet

Research Interests

Speech and language recognition; computational and statistical machine learning; large margin and kernel methods; online and iterative learning; speech coding; speech processing; medical ; spoken language understanding by humans.

Education

- 2002 - 2007 **Ph.D. in Computer Science**, The School of Computer Science and Engineering, The Hebrew University, Jerusalem, Israel.
Advisor: Yoram Singer
Dissertation: *Large Margin Algorithms for Discriminative Continuous Speech.*
- 2000 - 2002 **M.Sc. in Electrical Engineering**, Department of Electrical Engineering – Systems, Tel-Aviv University, Tel-Aviv, Israel.
Advisor: Dr. Dan Chazan, IBM Research Labs, Haifa.
Dissertation: *Stop consonant spotting in continuous speech.*
- 1991 - 1994 **B.Sc. in Electrical Engineering**, Cum laude
Department of Electrical Engineering – Systems,
Tel-Aviv University, Tel-Aviv, Israel.

Awards and Scholarships

- 2007 PASCAL Network of Excellence, Ph.D. Label Award.
- 2004-2006 Leibniz Student Fellowship.
- 2002 IBM Fellowship.
- 2001 Israel Defense Forces (IDF) – National award for outstanding research and development achievements.

Work Experience

- 2009 - present **Research Assistant Professor**, Toyota Technological Institute, University of Chicago, Chicago, IL.
- 2007 - 2009 **Research scientist (Postdoc)**, IDIAP Research Institute, Martigny, Switzerland.
- 2005 - 2007 **Consultant** to several speech and music recognition companies: NICE Systems Ltd., Neocraft Ltd., Semantrics, and Israeli Defense Forces (IDF).

- Summer 2006 **Internship** at IDIAP, Martigny, Switzerland.
Discriminative large margin approach for keyword spotting.
- Summer 2005 **Internship** at Google Inc., Mountain View, CA.
Scalability of speech recognition large-margin algorithms, specifically, phoneme sequence recognition and phonetic speech segmentation.
- 2002 - 2007 **Research Assistant**, School of Computer Science and Engineering,
The Hebrew University, Jerusalem, Israel.
- 1995 - 2002 **Head of speech processing research section**, The R&D Center,
Elite unit (8200), Israel Defense Forces (IDF).
Managing 2-3 researchers. Topics mainly concern speech coding, universal
speech encoder, information theoretic tools for speech coding and decoding.
- 1996 - 1998 **Senior signal processing engineer**, CTI² Ltd., Israel.
Research and development of speech coding techniques, speech streaming
and multi-rate signal processing.
- 1994 - 1995 **Signal processing engineer**, Comverse Inc., Israel.
Research and development of speech coding and speaker identification
algorithms.

Teaching Experience

- Guest Lecturer** Fall 2009, TTI-Chicago, for **Topics in Artificial Intelligence: Speech Technologies**.
- Guest Lecturer** Spring 2007, School of Engineering and Computer Science, The Hebrew
University of Jerusalem, for **Statistical Signal Processing**.
- Guest Lecturer** Fall 2007, School of Engineering and Computer Science, The Hebrew
University of Jerusalem, for **Digital Signal Processing**.
- Teaching Assistant** Spring 2006, School of Engineering, Bar-Ilan University, Ramat-Gan, for
Random Signals and Noise.
- Teaching Assistant** Fall 2005 and Fall 2006, School of Engineering and Computer Science, The
Hebrew University of Jerusalem, for **Digital Signal Processing**.
- Teaching Assistant** Fall 2004, School of Engineering and Computer Science, The Hebrew
University of Jerusalem, for **Control and Dynamic Systems**.
- Teaching Assistant** Fall 2003, School of Engineering and Computer Science, The Hebrew
University of Jerusalem, for **Information Theory**.
- Teaching Assistant** Spring 2002, Spring 2003, Spring 2004, and Spring 2005, School of
Engineering and Computer Science, The Hebrew University of Jerusalem,
for **Statistical Signal Processing**.

Selected Publications

Books and Chapters

Joseph Keshet and Samy Bengio, Eds., **Automatic Speech and Speaker Recognition: Large Margin
and Kernel Methods**, John Wiley & Sons, March 2009.

Joseph Keshet, Shai Shalev-Shwartz, Yoram Singer and Dan Chazan, **A Large Margin Algorithm for
Forced Alignment**, in *Automatic Speech and Speaker Recognition: Large Margin and Kernel Methods*, Joseph
Keshet and Samy Bengio, Eds., John Wiley & Sons, March 2009.

Joseph Keshet and Dan Chazan, **A Kernel Wrapper for Phoneme Sequence Recognition**, in *Automatic Speech and Speaker Recognition: Large Margin and Kernel Methods*, Joseph Keshet and Samy Bengio, Eds., John Wiley & Sons, March 2009.

Joseph Keshet, **A Proposal for a Kernel-based Algorithm for Large Vocabulary Continuous Speech Recognition**, in *Automatic Speech and Speaker Recognition: Large Margin and Kernel Methods*, Joseph Keshet and Samy Bengio, Eds., John Wiley & Sons, March 2009.

David Grangier, Joseph Keshet and Samy Bengio, **Discriminative Keyword Spotting**, in *Automatic Speech and Speaker Recognition: Large Margin and Kernel Methods*, Joseph Keshet and Samy Bengio, Eds., John Wiley & Sons, March 2009.

Journal Papers

Morgan Sonderegger and Joseph Keshet, **Automatic Discriminative Measurement of Voice Onset Time**, in submission to *Journal of the Acoustical Society of America*, 2011.

Andrew Cotter, Joseph Keshet, and Nathan Srebro, **Explicit Approximations of the Gaussian Kernel**, in submission to *Journal of Machine Learning Research*, 2011.
<http://arxiv.org/abs/1109.4603>

Francesco Orabona, Joseph Keshet and Barbara Caputo, **Bounded Kernel-Based Online Learning**, *Journal of Machine Learning Research*, 10(Nov): 2643–2666, 2009.

Joseph Keshet, David Grangier and Samy Bengio, **Discriminative Keyword Spotting**, *Speech Communication*, Volume 51, Issue 4, pp. 317-329, April 2009.

Joseph Keshet, Shai Shalev-Shwartz, Yoram Singer and Dan Chazan, **Large Margin Algorithm for Speech-to-Phoneme and Music-to-Score Alignment**, *IEEE Trans. on Audio, Speech and Language Processing*, 15(8), pp. 2373–2382, Nov. 2007.

Koby Crammer, Ofer Dekel, Joseph Keshet, Shai Shalev-Shwartz and Yoram Singer, **Online Passive-Aggressive Algorithms**, *Journal of Machine Learning Research*, 7(Mar): 551–585, 2006.

Conference papers

David McAllester and Joseph Keshet, **Generalization Bounds and Consistency for Latent Structural Probit and Ramp Loss**, Proc. of the 25th the Annual Conference on Neural Information Processing Systems (NIPS), 2011 (**full oral presentation**).

Joseph Keshet, Chih-Chieh Cheng, Mark Stoehr, and David McAllester, **Direct Error Rate Minimization of Hidden Markov Models**, Proc. of the 12th Annual Conference of the International Speech Communication Association (Interspeech), Florence, Italy, 2011.

Andrew Cotter, Nathan Srebro, and Joseph Keshet, **A GPU-Tailored Approach for Training Kernelized SVMs**, Proc. of the 17th ACM Conference on Knowledge Discovery and Data Mining (KDD), San Diego, CA, 2011.

Joseph Keshet, David McAllester, and Tamir Hazan, **PAC-Bayesian Approach for Minimization of Phoneme Error Rate**, Proc. of IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP), Prague, Czech Republic, 2011.

David McAllester, Tamir Hazan and Joseph Keshet, **Direct Loss Minimization for Structured Prediction**, Proc. of the 24th Annual Conference on Neural Information Processing Systems (NIPS), 2010.

Morgan Sonderegger and Joseph Keshet, **Automatic Discriminative Measurement of Voice Onset Time**, Proc. of the 12th Annual Conference of the International Speech Communication Association (INTERSPEECH), Chiba, Japan, 2010.

Martin Wöllmer, Florian Eyben, Joseph Keshet, Alex Graves, Björn Schuller, and Gerhard Rigoll, **Robust Discriminative Keyword Spotting for Emotionally Colored Spontaneous Speech using Bidirectional LSTM Networks**, Proc. of IEEE International Conference on Acoustic, Speech, and Signal Processing (ICASSP), Taipei, Taiwan, 2009.

Yves Grandvalet, Joseph Keshet, Alain Rakotomamonjy, and Stephane Canu, **Support Vector Machines with a Reject Option** Proc. of the 22nd Annual Conference on Neural Information Processing Systems (NIPS), 2008.

Francesco Orabona, Joseph Keshet and Barbara Caputo, **The Projectron: a Bounded Kernel-Based Perceptron** Proc. of the 25th International Conference on Machine Learning (ICML), 2008.

Joseph Keshet, David Grangier and Samy Bengio, **Discriminative Keyword Spotting**, Proc. of the 3rd Non-Linear Speech Processing Workshop (NOLISP), Paris, 2007.

Joseph Keshet, Shai Shalev-Shwartz, Samy Bengio, Yoram Singer and Dan Chazan, **Large Margin Phoneme Sequence Recognition**, Proc. of the 9th International Conference on Spoken Language Processing (Interspeech), Pittsburgh, PA, 2006.

Joseph Keshet, Shai Shalev-Shwartz, Yoram Singer and Dan Chazan, **Phoneme Alignment Based on Discriminative Learning**, Proc. of the 9th European Conference on Speech Communication and Technology (Interspeech), Lisbon, 2005.

Shai Shalev-Shwartz, Joseph Keshet and Yoram Singer, **Learning to Align Polyphonic Music**, Proc. of the 5th International Conference on Music Information Retrieval (ISMIR), Barcelona, Spain, 2004.

Ofer Dekel, Joseph Keshet and Yoram Singer, Workshop on Multimodal Interaction and Related **An Online Algorithm for Hierarchical Phoneme Classification**, Machine Learning Algorithms (MLMI), Martigny, Switzerland, 2004, LNCS, Vol. 3361/2005, p. 146, Springer-Verlag.

Ofer Dekel, Joseph Keshet and Yoram Singer, **Large Margin Hierarchical Classification**, Proc. of the 21st International Conference on Machine Learning (ICML), Banff, Canada, 2004.

Koby Crammer, Joseph Keshet and Yoram Singer, **Kernel Design using Boosting**, in Proc. of 16th Annual Conference on Neural Information Processing Systems (NIPS), 2002.

Joseph Keshet, Dan Chazan and Ben-Zion Bobrovsky, **Plosive Spotting with Large Margin Classifiers**, in Proc. of 7th European Conference on Speech Communication and Technology (Eurospeech), 2001.

Selected Invited Talks

Colloquium, School of Engineering, University of Southern California, “Loss Minimization for VOT Measurement, Phoneme Alignment, and Phoneme Recognition”, Oct 2011.

Colloquium, CS, U. of Washington, “Direct Loss Minimization for Structural Prediction with Applications to Speech Recognition”, Oct 2011.

The LTI colloquium, CS, Carnegie-Mellon University, Pittsburgh, PA, "Loss Minimization for VOT Measurement, Phoneme Alignment, and Phoneme Recognition", Sept 2011.

Colloquium, Dep. of Science and Engineering, Oregon Health and Science University, "Loss Minimization for VOT Measurement, Phoneme Alignment, and Phoneme Recognition", April 2011.

Invited talk, Workshop on New Tools and Methods for Very-Large-Scale Phonetics Research, University of Pennsylvania, "Accurate Discriminative Forced Alignment", Jan 2011.

Workshop on New Tools and Methods for Very-Large-Scale Phonetics Research, University of Pennsylvania, "A new tool for automatic VOT measurement", Jan 2011.

Workshop on New Tools and Methods for Very-Large-Scale Phonetics Research, University of Pennsylvania, "Automatic keyword spotting for continuous speech", Jan 2011.

Google Research Labs, Mountain View, CA, "Direct Loss Minimization for Structural Prediction with Applications to Speech Recognition", Jan 2011.

IBM colloquium, IBM Watson Research Center, Yorktown, NY, "Direct Loss Minimization for Structural Prediction with Applications to Speech Recognition", Feb 2011.

EESIP Seminar, Columbia University, "Direct Loss Minimization for Structural Prediction with Applications to Speech Recognition", Feb 2011.

CS Colloquium, University of California, San Diego, CA, "Direct Loss Minimization for Structural Prediction with Applications to Speech Recognition", Nov 2010.

NLP Colloquium (Dan Roth's group), University of Illinois at Urbana-Champaign, IL, "Direct Loss Minimization for Structural Prediction with Applications to Speech Recognition", Nov 2010.

CS Colloquium, Hebrew University, Jerusalem, Israel, "Speech Recognition with Structured Labeling", Sept 2010.

Illinois Speech Day, TTI, Chicago, IL, "Direct Loss Minimization for Forced Alignment", May 2010.

FBK and Trento University, Trento, Italy, "Discriminative Spoken Keyword Detection", Dec 2009

EESIP Seminar, Columbia University, New-York, NY, "Discriminative Spoken Keyword Detection", Dec 2009

IBM Zurich Labs, "Discriminative Keyword Spotting", August 2008.

PASCAL Symposium, Bled, Slovenia, "Pump Priming Project: Large Margin Methods for Speech Recognition", January 2008.

Speech Technologies Day at IBM Haifa, Israel, "Discriminative Keyword Spotting", July 2008.

Learning Club, School of Computer Science, The Hebrew University, Jerusalem, Israel, "Large Margin Phoneme Sequence Recognition", May 2006.

Colloquium, ICSI, Berkeley, CA, "Phoneme Alignment Based on Discriminative Learning", August 2005.

The Natural Language and Speech Processing Colloquium, NLP Group, Stanford, Palo-Alto, CA, "Phoneme Alignment Based on Discriminative Learning", August 2005.

Google Research Labs, Mountain View, CA, "Phoneme Alignment Based on Discriminative Learning", July 2005.

Professional Services

Journal Refereeing

Speech Communication, IEEE Trans. on Speech and Audio Processing, IEEE Journal of Selected Topics in Signal Processing, EURASIP Journal on Audio, Speech, and Music Processing, Journal of Machine Learning Research (JMLR), Pattern Recognition, Pattern Recognition Letters,

Conference Refereeing

Neural Information Processing Systems (NIPS), ICML

Appointments

Founder and organizer the first symposium on Machine Learning for Speech and Language Processing, Bellevue, WA, June 27, 2011.

<http://www.ttic.edu/sigml/symposium2011>

Founder and Liaison of the Machine Learning for Speech and Language Processing Special Interest Group (SIGML) of International Speech Communication Association (ISCA), 2011.

Organizer of Illinois Speech Day 2011.

http://ttic.uchicago.edu/~jkeshet/Illinois_Speech_Day_2011.html

Organizer of the Learning Club at the Hebrew University, 2006-2007.

Student advisory committee of the International Speech Communication Association, (ISCA), 2006-2007.